A Risk Ratio Comparison of l_0 and l_1 Penalized Regression

Kory D. Johnson[1], Dongyu Lin[2], Lyle H. Ungar[3], Dean P. Foster[1], and Robert A. Stine[1]

[1] *The Wharton School, University of Pennsylvania*
[2] *AT&T Labs*
[3] *The School of Engineering and Applied Science, University of Pennsylvania*

October 22, 2015

arXiv:1510.06319v1 [math.ST] 21 Oct 2015

Abstract

There has been an explosion of interest in using l_1-regularization in place of l_0-regularization for feature selection. We present theoretical results showing that while l_1-penalized linear regression never outperforms l_0-regularization by more than a constant factor, in some cases using an l_1 penalty is infinitely worse than using an l_0 penalty. We also show that the "optimal" l_1 solutions are often inferior to l_0 solutions found using stepwise regression.

We also compare algorithms for solving these two problems and show that although solutions can be found efficiently for the l_1 problem, the "optimal" l_1 solutions are often inferior to l_0 solutions found using greedy classic stepwise regression. Furthermore, we show that solutions obtained by solving the convex l_1 problem can be improved by selecting the best of the l_1 models (for different regularization penalties) by using an l_0 criterion. In other words, an approximate solution to the right problem can be better than the exact solution to the wrong problem.

Keywords: Variable Selection, Streaming Feature Selection, Regularization, Stepwise Regression, Submodularity

1 Introduction

In the past decade, a rich literature has been developed using l_1-regularization for linear regression including Lasso (Tibshirani, 1996), LARS (Efron et al., 2004), fused lasso (Tibshirani et al., 2005), elastic net (Zou and Hastie, 2005), grouped lasso (Yuan and Lin, 2006), adaptive lasso (Zou, 2006), and relaxed lasso (Meinshausen, 2007). These methods, like the l_0-penalized regression methods which preceded them (Akaike, 1974; Schwarz, 1978; Foster and George, 1994), address variable selection problems in which there is a large set of potential features, only a few of which are likely to be helpful. This type of sparsity is common in machine learning tasks, such as predicting disease based on thousands of genes, or predicting the topic of a document based on the occurrences of hundreds of thousands of words.

l_1-regularization is popular because, unlike the l_0 regularization historically used for feature selection in regression problems, the l_1 penalty gives rise to a convex problem that can be solved efficiently using convex optimization methods. l_1 methods have given reasonable results on a number of data sets, but there has been no careful analysis of how they perform when compared to l_0 methods. This paper provides a formal analysis of the two methods, and shows that l_1 can give arbitrarily worse models. We offer some intuition as to why this is the case – l_1 shrinks coefficients too much and does not zero out enough of them – and suggest how to use an l_0 penalty with l_1 optimization.

We study the problem of selecting predictive features from a large feature space. We assume the classic normal linear model

$$\mathbf{y} = \mathbf{X}\beta + \epsilon \qquad\qquad \epsilon \sim N_n(\mathbf{0}, \sigma^2\mathbf{I}_n)$$

with n observations $\mathbf{y} = (y_1, \ldots, y_n)'$ and p features $\mathbf{x}_1, \ldots, \mathbf{x}_p$, where $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_p)$ is an $n \times p$ "design matrix" of features, $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)'$ is the coefficient parameters, and error $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 I_n)$.

We expect most of the elements of β to be 0. Hence, generating good predictions requires identifying the small subset of predictive features. This standard linear model proliferates the statistics and machine learning literature. In modern applications, p can approach millions, making the selection of an

appropriate subset of these features essential for prediction. The size and scope of these problems raise concerns about both the speed and statistical robustness of the selection procedure. Namely, it must be fast enough to be computationally feasible and must find signal without over-fitting the data.

The traditional statistical approach to this problem, namely, the l_0 regularization problem, finds an estimator that minimizes the l_0 penalized sum of squared errors,

$$\arg\min_{\boldsymbol{\beta}} \left\{ \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 + \lambda_0 \|\boldsymbol{\beta}\|_{l_0} \right\}, \tag{1}$$

where $\|\boldsymbol{\beta}\|_{l_0} = \sum_{i=1}^{p} I_{\{\beta_i \neq 0\}}$ counts the number of nonzero coefficients. However, this problem is NP-hard (Natarajan, 1995). A tractable problem relaxes the l_0 penalty to the l_1 norm, $\|\boldsymbol{\beta}\|_{l_1} = \sum_{i=1}^{p} |\beta_i|$, and seeks

$$\arg\min_{\boldsymbol{\beta}} \left\{ \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 + \lambda_1 \|\boldsymbol{\beta}\|_{l_1} \right\}, \tag{2}$$

This is known as the l_1-regularization problem (Tibshirani, 1996), which is convex. This problem can be solved efficiently using a variety of methods (Tibshirani, 1996; Efron et al., 2004; Candes and Tao, 2007).

We assess our models using the predictive risk function (3)

$$R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}) = \mathbb{E}_{\boldsymbol{\beta}} \|\hat{\mathbf{y}} - \mathbb{E}(\mathbf{y}|\mathbf{X})\|_2^2 = \mathbb{E}_{\boldsymbol{\beta}} \|\mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}\|_2^2. \tag{3}$$

We are interested in the ratios of the risks of the estimates provided by these two criteria. Unlike risk functions, predictive risk measures the relevant prediction error on future observations, ignoring irreducible variance. Smaller risks imply better expected prediction performance. It is an ideal metric to analyze testing error or out-of-sample errors when the parameter distribution is assumed to be known. Recent literature has focused on selection consistency: whether or not the true variable can be identified in the limit. However, in real application, due to ubiquitous multicollinearity, predictors are hard to separate as "true" and "false". Here, we focus on predictive accuracy and advocate the concept of predictive risk.

Our first result in this paper, given below as Theorems 1 and 2 in Section 2, is that l_0 estimates provide more accurate predictions than l_1 estimates do, in the sense of minimax risk ratios. This is illustrated in Figure 1. Proofs of these theorems are in Appendix A.

- $\inf_{\gamma_0} \sup_{\boldsymbol{\beta}} \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))}$ is bounded by a small constant; furthermore, it is close to one for most γ_1s, especially for large γ_1s, which are mostly used in sparse systems.
- $\inf_{\gamma_1} \sup_{\boldsymbol{\beta}} \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))}$ tends to infinity quadratically; in an extremely sparse system, the l_1 estimate may perform arbitrarily badly.
- $R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))$ is more likely to have a larger risk than $R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))$ does.

A detailed discussion on the risk ratios will be presented in Section 2, along with a discussion of other advantages of l_0 regularization. Our other comparative results include showing that applying the l_0 criterion on an l_1 subset searching path can find the best performing model (Section 2.3.1) and running stepwise regression and Lasso on a reduced NP hard example shows that stepwise regression gives better solutions (Section 2.3.2).

We compare l_0 vs. l_1 penalties under three assumptions about the structure of the feature matrix X: independence, incoherence (near independence) and NP-hard. For independence, we find provide the theoretical results mentioned above. For near independence, we find that l_1 penalized regression followed by l_0 outperforms l_1 selection. For the NP-hard case, we find that if one could do the search, then the risk ratio could be arbitrarily bad for l_1 relative to l_0.

2 Risk Ratio Comparison

We assess our models using the predictive risk function (3)

$$R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}) = \mathbb{E}_{\boldsymbol{\beta}} \|\hat{\mathbf{y}} - \mathbb{E}(\mathbf{y}|\mathbf{X})\|_2^2 = \mathbb{E}_{\boldsymbol{\beta}} \|\mathbf{X}\hat{\boldsymbol{\beta}} - \mathbf{X}\boldsymbol{\beta}\|_2^2. \tag{4}$$

This is the relevant component after decomposing the expected squared error loss from predicting a new observation. This is clear from the following standard decomposition. For increased generality, let $\mathbb{E}[y] = \eta]$ and $\mathbf{H}_{\mathbf{X}}$ be the projection onto the column space of $\mathbf{X}$. Then

$$\begin{aligned} \mathbb{E}\|y^* - \mathbf{X}\hat{\beta}\|^2 &= \mathbb{E}\|y^* - \eta\|^2 + \mathbb{E}\|\eta - \mathbf{H}_{\mathbf{X}}\eta\|^2 + \mathbb{E}\|\mathbf{H}_{\mathbf{X}}\eta - \mathbf{X}\hat{\beta}\|^2 \\ &= \underbrace{n\sigma^2}_{\text{common error}} + \underbrace{\|(\mathbf{I} - \mathbf{H}_{\mathbf{X}})\eta\|^2}_{\text{wrong } \mathbf{X}} + \underbrace{\mathbb{E}\|\mathbf{H}_{\mathbf{X}}\eta - \mathbf{X}\hat{\beta}\|^2}_{\text{predictive risk}} \end{aligned}$$

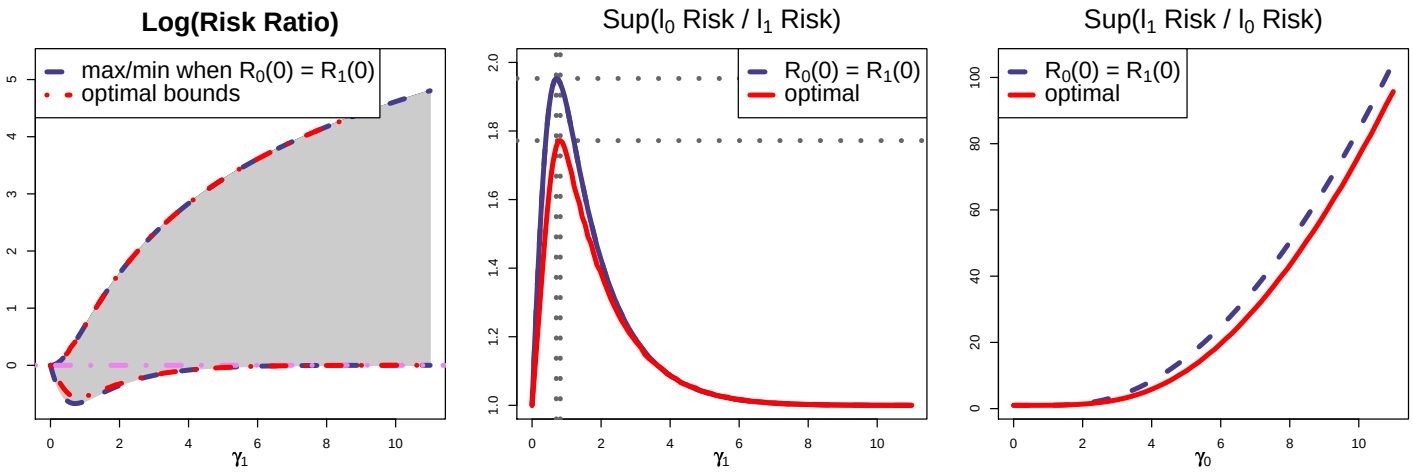

Figure 1: **Left**: The gray area shows the feasible region for the risk ratios–the log risk-ratio is above zero when l_0 produces a better fit. The graph shows that most of the time l_0 is better. The actual estimators being compared are those that have the same risk at $\beta = 0$, i.e., $R(0, \hat{\beta}_{l_0}(\gamma_0)) = R(0, \hat{\beta}_{l_1}(\gamma_1))$. **Middle**: This graph traces out the bottom envelope of the left hand graph (but takes the reciprocal risk ratio and no longer uses the logarithm scale). The dashed blue line displays $\sup_\beta R(\beta, \hat{\beta}_{l_0}(\gamma_0))/R(\beta, \hat{\beta}_{l_1}(\gamma_1))$ for γ_0 calibrated to have the same risk at zero as γ_1. This maximum ratio tends to 1 when $\gamma_1 \to 0$ (the saturated case) or ∞ (the sparse case). With an optimal choice of γ_0, $\inf_{\gamma_0} \sup_\beta R(\beta, \hat{\beta}_{l_0}(\gamma_0))/R(\beta, \hat{\beta}_{l_1}(\gamma_1))$ (solid red line) behaves similarly. Specifically, the supremum over γ_1 is bounded by 1.8. **Right**: This graph traces out the upper envelopes of the left hand graph on a normal scale. When $\gamma_0 \to \infty$, $\sup_\beta R(\beta, \hat{\beta}_{l_1}(\gamma_1))/R(\beta, \hat{\beta}_{l_0}(\gamma_0))$ tends to ∞, for both γ_1 that is calibrated at $\beta = 0$ and that minimizes the maximum risk ratio.

The first term, common error, is unavoidable, regardless of the method being used. All methods we consider, namely linear methods based on $\mathbf{X}$, suffer the error from incorrect $\mathbf{X}$. Since $\mathbf{X}$ is given, it is more instructive to consider the projection of η onto the column space of $\mathbf{X}$, defining $\mathbf{X}\beta = \mathbf{H}_\mathbf{X}\eta$. Ignoring these two forms of error, leaves the predictive risk function (3).

Predictive risk has guided selection procedures such as Mallow's C_p and RIC. The former results from an unbiased estimate of the predictive risk, while the later provides minimax control of the risk in during model selection. We maintain this minimax viewpoint and show that in terms of the *removable* variation in prediction, l_0 performs better than l_1.

2.1 l_0 solutions give more accurate predictions.

Suppose that $\hat{\boldsymbol{\beta}}$ is an estimator of $\boldsymbol{\beta}$. For this section, we assume $\mathbf{X}$ is orthogonal. (For example, wavelets, Fourier transforms, and PCA all are orthogonal). The l_0 problem (1) can then be solved by simply picking those predictors with least squares estimates $|\hat{\beta}_i| > \gamma$, where the choice of γ depends on the penalty λ_0 in (1). It was shown (Donoho and Johnstone, 1994; Foster and George, 1994) that $\lambda_0 = 2\sigma^2 \log p$ is optimal in the sense that it asymptotically minimizes the maximum predictive risk inflation due to selection.

Let

$$\hat{\boldsymbol{\beta}}_{l_0}(\gamma_0) = \left(\hat{\beta}_1 I_{\{|\hat{\beta}_1|>\gamma_0\}}, \ldots, \hat{\beta}_p I_{\{|\hat{\beta}_p|>\gamma_0\}}\right)' \tag{5}$$

be the l_0 estimator that solves (1), and let the l_1 solution to (2) be

$$\hat{\boldsymbol{\beta}}_{l_1}(\gamma_1) = \left(\text{sign}(\hat{\beta}_1)(|\hat{\beta}_1| - \gamma_1)_+, \ldots, \text{sign}(\hat{\beta}_p)(|\hat{\beta}_p| - \gamma_1)_+\right)', \tag{6}$$

where the $\hat{\beta}_i$'s are the least squares estimates.

We are interested in the ratios of the risks of these two estimates,

$$\frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))} \quad \text{and} \quad \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))}.$$

I.e., we want to know how the risk is inflated when another criterion is used. The smaller the risk ratio, the less risky (and hence better) the numerator estimate is, compared to the denominator estimate. Specifically, a risk ratio less than one implies that the top estimate is better than the bottom estimate.

Formally, we have the following theorems, whose proofs are given in Appendix A:

Theorem 1. *There exists a constant C_1 such that for any $\gamma_0 \geq 0$,*

$$\inf_{\gamma_1} \sup_{\boldsymbol{\beta}} \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))} \geq C_1 + \gamma_0.$$

I.e., given γ_0, for any γ_1, there exist $\boldsymbol{\beta}$'s such that the ratio becomes extremely large.
Contrast this with the protection provided by l_0:

Theorem 2. *There exists a constant $C_2 > 0$ such that for any $\gamma_1 \geq 0$,*

$$\inf_{\gamma_0} \sup_{\boldsymbol{\beta}} \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))} \leq 1 + C_2 \gamma_1^{-1}.$$

I.e., for any γ_1, we can pick the l_0 cutoff so that we perform almost as good as l_1, even in the worst case.

The above theorems can definitely be strengthened, as demonstrated by the bounds shown in Figure 1, but at the cost of complicating the proofs. We conjecture that there exist constants $r > 1$, and $C_3, C_4, C_5 > 0$, such that

$$\inf_{\gamma_1} \sup_{\boldsymbol{\beta}} \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))} \geq 1 + C_3 \gamma_0^r, \tag{7}$$

$$\inf_{\gamma_0} \sup_{\boldsymbol{\beta}} \frac{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}(\gamma_0))}{R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}(\gamma_1))} \leq 1 + C_4 \gamma_1 e^{-C_5 \gamma_1}. \tag{8}$$

These theorems suggest that for any γ_1 chosen by the algorithm, we can always adapt γ_0 such that $\hat{\boldsymbol{\beta}}_{l_0}(\gamma_0)$ outperforms $\hat{\boldsymbol{\beta}}_{l_1}(\gamma_1)$ most of the time and loses out a little for some $\boldsymbol{\beta}$'s; but for any γ_0 chosen, no γ_1 can perform consistently well on all $\boldsymbol{\beta}$'s.

Because of the additivity of risk functions, (see appendix equations (13) and (14)), due to the orthogonality assumption, we focus on the individual behavior of β_i for each single feature. Also the risk functions are symmetric on β, so only the cases of $\beta_i \geq 0$ will be displayed. Figure 2.1 illustrates that given γ_1, we can pick a γ_0, s.t. the risk ratio is below 1 for most β except around $(\gamma_0 + \gamma_1)/2$, yet this ratio does not exceed one by more than a small factor, even for the worst case.

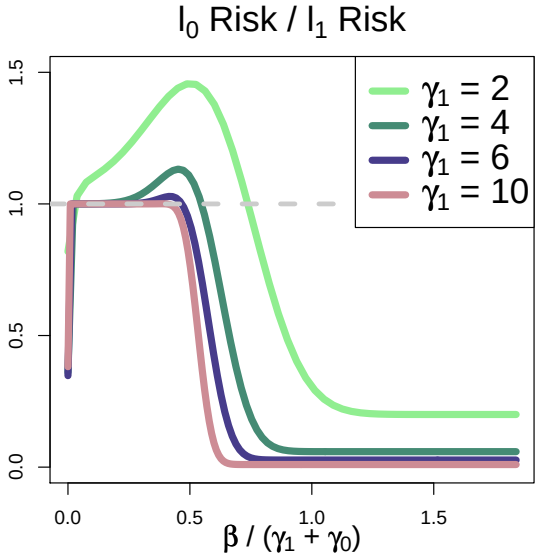

Figure 2: For each γ_1, we let $\gamma_0 = \gamma_1 + 4\log(\gamma_1)/\gamma_1$. This choice of γ_0 makes the risk ratios small at $\beta \approx 0$ and $\beta \geq \gamma_0$, only inflated around $\beta/(\gamma_0 + \gamma_1) = 1/2$, albeit very little especially when γ_1 is large enough.

The intuition as to why l_0 fares better than l_1 in the risk ratio results is that l_1 must make a "devil's choice" between shrinking the coefficients too much or putting in too many spurious features. l_0 penalized regression avoids this problem.

2.2 l_1 shrinks coefficients too much

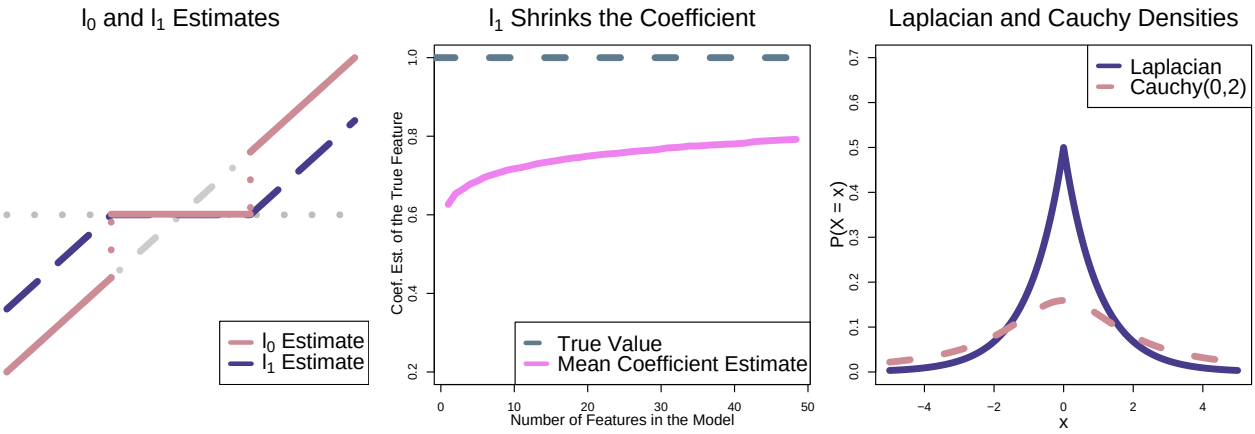

Figure 3: **Left**: The l_0 estimate keeps the least squares value after the cutting point, but the l_1 estimate always shrinks the least squares estimate by a fixed amount. **Middle**: the model we simulate has only one true feature with true $\beta = 1$ and a thousand spurious features. We compute the average Lasso estimate of β for a fixed number of features included in the model (as an index of the l_1 penalty) from several different trials. $\bar{\hat{\beta}}_{l_1}$ is always shrunk by at least 20% in this experiment. **Right**: The Cauchy density has heavier tails than the Laplacian density does. Thus, a Laplacian prior tends to shrink large values of β's.

From a frequentist's point of view, the l_1 estimator (6) shrinks the coefficients and thus is biased (Figure 3). In practice, $\hat{\beta}_{l_1}$ can be substantially shrunk towards zero when the system is sparse, as shown in the middle panel of Figure 3.

From a Bayesian's perspective, the l_1 penalty is equivalent to putting a Laplacian prior on β (Tibshirani, 1996; Efron et al., 2004), while the l_0 penalty can be approximated by Cauchy priors (Johnstone and Silverman, 2005; Foster and Stine, 2005). The right panel of Figure 3 shows that the Cauchy distribution has a much heavier tail than the Laplacian distribution does. This implies that when the true β is far away from 0, the l_1 penalty will substantially shrink the estimate toward zero.

The bias caused by the shrinkage increases the predictive risk proportionally to the squared amount of the shrinkage. The sparser the problem is, the greater the shrinkage is, thus the larger the risk is. These results show that in theory the l_0 estimate has a lower risk and provides a more accurate solution. Empirically, stepwise regression performs well in large data sets, where a sparse solution is particularly preferred (George and Foster, 2000; Foster and Stine, 2004; Zhou et al., 2006).

2.3 Simulations for Risk Ratio/flaws with l1

2.3.1 l_1 optimization using an l_0 criterion

We can make use of the LARS algorithm to generate a set of candidate solutions and then use the l_0 criterion to find the best of the solutions along the regularization path. We evaluated this method as follows. We simulated $\mathbf{y}$ from a thousand features, only 4 of which have nonzero contributions, plus a random noise distributed as $N(0, 1)$. Both the training set and the test set have size $n = 100$. We apply the Lasso algorithm implemented by LARS on this synthetic data set. For each step on the regularization path, this algorithm selects a subset $\mathcal{C} \subset \{1, \ldots, 1000\}$ of features that are included in the model. We then adopt a modified RIC criterion suggested in George and Foster (2000):

$$\|\mathbf{y} - \mathbf{X}_{\mathcal{C}}\hat{\boldsymbol{\beta}}_{\mathcal{C}}\|_2^2 + \sum_{q=1}^{|\mathcal{C}|} 2\log(p/q)\sigma^2 \tag{9}$$

to find an optimal $\mathcal{C}$. The crucial part here is that the coefficient estimate $\hat{\boldsymbol{\beta}}_{\mathcal{C}}$ being used in (9) is the least squares estimate of the true $\boldsymbol{\beta}$ obtained by fitting $\mathbf{y}$ on $\mathbf{X}_{\mathcal{C}} = (\mathbf{x}_j)_{j\in\mathcal{C}}$, and not the Lasso estimate $\hat{\boldsymbol{\beta}}_{l_1}$ provided by the algorithm. We also use this least squares estimate in out-of-sample calculations.

We compare two cases: the $\mathbf{x}_j$'s are generated independently of each other, meaning that $X'X$ is diagonal, and the $\mathbf{x}_j$'s are generated with a pairwise correlation $\rho = 0.64$. As shown in Figure 4, in the

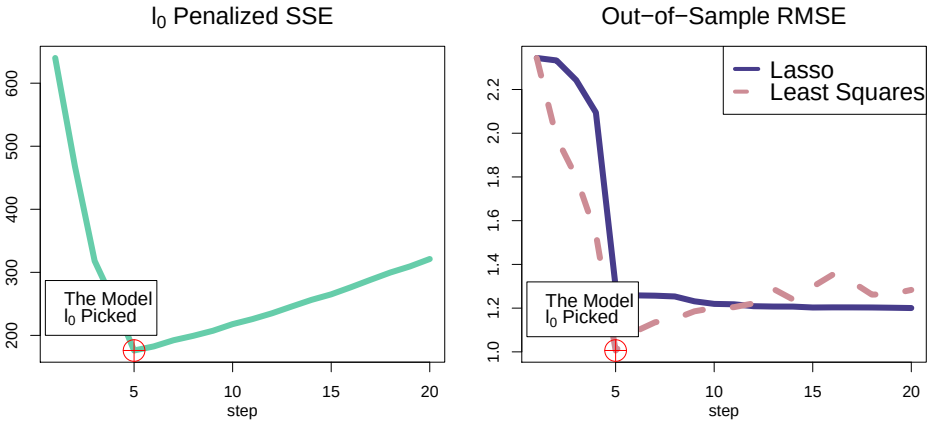

Figure 4: l_0 penalties help finding the best model (independent predictors case). $\mathbf{y}$ is simulated from one thousand features, only four of which have nonzero contributions, plus an $N(0,1)$ error. Both the training set and the test set have sizes $n = 100$. Each step in the LARS algorithm gives a set of features with nonzero coefficient estimates. We compute the least squares (LS) estimates on this subset and the modified RIC criterion (9) on the training set. We also compare the out-of-sample root mean squared errors using the LS estimates and the Lasso estimates on this LARS path. The features are independently generated. The model that minimizes the l_0 penalized error has exactly four variables in it. It also outperforms any of the l_1 models out-of-sample on this data set.

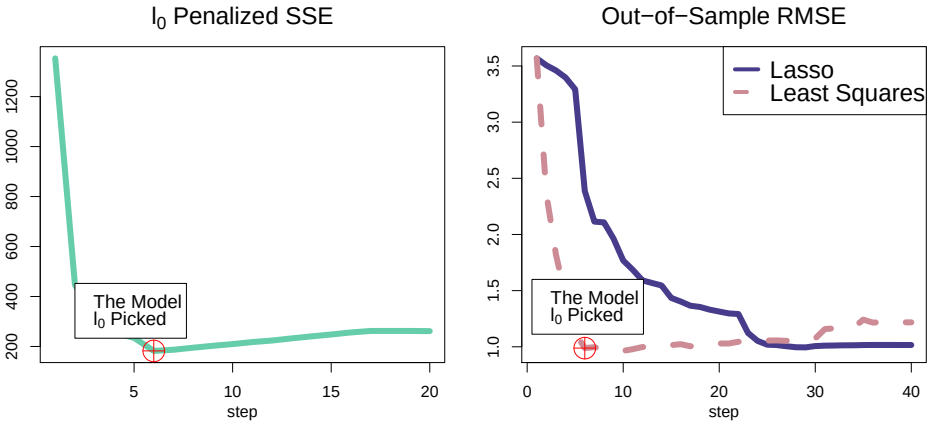

Figure 5: l_0 penalties help finding the best model (correlated predictors case). The setup is exactly the same as in Figure 4 except that each pair of features has a correlation $\rho = 0.64$. In this case, the optimal model under the modified RIC criterion has a slightly better RMSE than the best l_1 model. The Lasso out-of-sample RMSE is typically minimized when the model has included more than 50 features.

independent feature case, the model picked by the modified RIC criterion always outperforms any Lasso model on the test set. In the case with correlated predictors (Figure 5), there is little difference between the out-of-sample accuracies of the l_0-picked model and the best Lasso model in this case, but Lasso adds around 50 more spurious variables.

Thus, by combining the computational efficiency of an l_1 algorithm and the sparsity guaranteed by the l_0 penalization, we can easily select an accurate model without cross validation.

2.3.2 l_0 and NP-hardness

The l_0 problem is NP-hard and hence, at least in theory, intractable. (In practice, of course, people often use approximate solutions to problems that in the worst case can be NP-hard.) One of the attractions of l_1-regularization is that it is convex, hence solvable in polynomial time. In this section, we compare how the two approaches fare on a known NP-hard regression problem. We start with a simple constructive proof that the risk ratio for l_1 to l_0 can be arbitrarily bad. Construct data as follows. Pick a large number of independent features z_j. Construct new features $x_1 = z_1 + \epsilon z_2$ and $x_2 = z_1 - \epsilon z_2$ and. Let $y = (z_1 + z_2)/2$ plus noise. Then the correct model is $y = x_2/\epsilon$. Include the rest of the features $z_j, j > 2$ as spurious features.

In Natarajan (1995) the known NP hard problem of "the exact cover of 3-sets" was reduced to the best subset selection problem as below: $\mathbf{y} = \mathbf{1}_n$, X is an $n \times p$ binary matrix with each column having three nonzero elements: $\|\mathbf{x}_i\|_0 = 3$, $\boldsymbol{\beta}$ is a $p \times 1$ vector, $\varepsilon > 0$ and we want to solve

$$\min_{\boldsymbol{\beta}} \|\boldsymbol{\beta}\|_0, \quad \text{s.t.} \quad \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2 < \varepsilon. \tag{10}$$

Note that if there *is* a solution to this problem, the number of features being chosen should be $n/3$.

We then ask which method comes closer to solving this problem: a greedy approximation to the l_0 problem or an exact solution to the l_1 problem. To this end, we applied Lasso and forward stepwise regression on various n's. For small n's, we took full collections of the three subsets, i.e., p equals n choose 3; for larger n's, we took $p = 10 \cdot n$. Table 1 and 2 list the number of subsets included in the model. Forward stepwise regression always finds fewer subsets, and hence a better solution, than Lasso.

Method	$n = 9$	$n = 12$	$n = 15$	$n = 18$	$n = 21$	$n = 24$	$n = 27$	$n = 30$
Lasso	6	10	11	17	19	21	22	29
Stepwise	3	4	5	6	7	8	9	10

Table 1: The number of subsets chosen by Lasso and by forward stepwise regression with $\varepsilon = 1/4$. All 3-subsets were considered, i.e., $p = \binom{n}{3}$. Forward stepwise regression always has the fewest possible number of subsets, namely, $n/3$.

Method	$n = 99$	$n = 240$	$n = 540$	$n = 990$	$n = 1500$
Lasso	93	219	504	812	1372
	(2×10^{-23})	(9×10^{-23})	(9×10^{-15})	(6×10^{-20})	(2×10^{-20})
Stepwise	40	96	223	364	595
	(1×10^{-28})	(6×10^{-27})	(3×10^{-26})	(6×10^{-25})	(1×10^{-25})

Table 2: The number of subsets chosen by Lasso and by forward stepwise regression with $\varepsilon = 1/4$. $p = 10 \cdot n$ 3-subsets were randomly chosen to be the predictors. Forward stepwise regression always chooses a sparser solution in the sense that it chooses fewer number of subsets. Numbers in parentheses are the sum of squared errors when the algorithms terminated.

All of our experiments on both synthetic and real data sets show that greedy search algorithms, such as stepwise regression, aimed at minimizing l_0-regularized error provide sparser results. This is because l_0 penalizes the sparsity directly, while l_1 does not. It is easy to construct an example where l_1 will pick a solution with a smaller l_1 norm but with a less sparse solution (Candes et al., 2007).

3 Conclusion

In many statistical contexts, the l_0 regularization criterion is superior to that of l_1 regularization; l_0 generally provides a more accurate solution and controls the false discovery rate better. l_1 can give

arbitrarily worse predictive accuracy than l_0, since l_1 regularization tends to shrink coefficients too much to include many spurious features. Computationally, l_1 appears to be more attractive; convex programming makes the computation feasible and efficient. In practice, however, approximate solutions to the l_0 problem are often better than than exact solutions to the l_1 problem. The best properties of the two methods can be combined. Superior results were obtained by using convex optimization of the l_1 problem to generate a set of candidate models (the regularization path generated by LARS), and then selecting the best model by minimizing the l_0-penalized training error.

Appendices

A Risk Ratio Proofs

We will drop the γ's when the situation is clear, and denote $\hat{\beta}_{l_0}(\gamma_0)$ as $\hat{\beta}_{l_0}$ and $\hat{\beta}_{l_1}(\gamma_1)$ as $\hat{\beta}_{l_1}$ for simplicity.

Without loss of generality, we assume $\mathbf{X}'\mathbf{X} = I$ and $\sigma = 1$. The l_0 risk can be written as

$$\begin{aligned} R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_0}) &= \mathbb{E}_{\boldsymbol{\beta}} \|\mathbf{X}\boldsymbol{\beta} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2 = \mathbb{E}_{\boldsymbol{\beta}} \sum_{i=1}^{p} \|\mathbf{x}_i\|^2 (\beta_i - \hat{\beta}_i)^2 \\ &= \mathbb{E}_{\boldsymbol{\beta}} \sum_{i=1}^{p} \left(\left(\frac{\mathbf{x}_i'\boldsymbol{\varepsilon}}{\|\mathbf{x}_i\|} \right)^2 I_{\{|\hat{\beta}_i| > \gamma\}} + (\|\mathbf{x}_i\|\beta_i)^2 I_{\{|\hat{\beta}_i| \le \gamma\}} \right) \\ &= \sum_{i=1}^{p} \left\{ \sigma^2 \mathbb{E}_\beta \left[Z_i^2 I_{\{|\beta_i + \sigma Z_i| > \gamma\}} \right] + (\|\mathbf{x}_i\|\beta_i)^2 P(|\beta_i + \sigma Z_i| \le \gamma) \right\}, \end{aligned} \tag{11}$$

where $Z_i = \mathbf{x}_i'\boldsymbol{\varepsilon}/\sigma\|\mathbf{x}_i\| \sim N(0,1)$, $i = 1, \ldots, p$.

Similarly, the l_1 risk can be written as

$$\begin{aligned} R(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}}_{l_1}) = \mathbb{E}_{\boldsymbol{\beta}} \sum_{i=1}^{p} \Bigg(& \left(\frac{\mathbf{x}_i'\boldsymbol{\varepsilon}}{\|\mathbf{x}_i\|} - \tilde{\gamma} \right)^2 I_{\{\hat{\beta}_i > \tilde{\gamma}\}} + \left(\frac{\mathbf{x}_i'\boldsymbol{\varepsilon}}{\|\mathbf{x}_i\|} + \tilde{\gamma} \right)^2 I_{\{\hat{\beta}_i < -\tilde{\gamma}\}} \\ & + (\|\mathbf{x}_i\|\beta_i)^2 I_{\{|\hat{\beta}_i| \le \tilde{\gamma}\}} \Bigg) \\ = \sum_{i=1}^{p} \Big\{ & \mathbb{E}_\beta \left[(\sigma Z_i - \tilde{\gamma})^2 I_{\{\beta_i + \sigma Z_i > \tilde{\gamma}\}} + (\sigma Z_i + \tilde{\gamma})^2 I_{\{\beta_i + \sigma Z_i < -\tilde{\gamma}\}} \right] \\ & + (\|\mathbf{x}_i\|\beta_i)^2 P(|\beta_i + \sigma Z_i| \le \gamma) \Big\} \end{aligned} \tag{12}$$

Specifically, we consider the case when $p = 1$. Let $\Phi(z) = P(Z \le z)$ and $\tilde{\Phi}(z) = P(Z > z)$ be the lower and upper tail probabilities of a standard normal distribution and the two risk functions can be explicitly written as

$$\begin{aligned} R(\beta, \hat{\beta}_{l_0}) &= \int_{\gamma_0 - \beta}^{\infty} z^2 \phi(z)\, \mathrm{d}z + \int_{-\infty}^{-\gamma_0 - \beta} z^2 \phi(z)\, \mathrm{d}z + \beta^2 \left[\Phi(\gamma_0 - \beta) - \tilde{\Phi}(\gamma_0 + \beta) \right] \\ &= (\gamma_0 - \beta)\phi(\gamma_0 - \beta) + (\gamma_0 + \beta)\phi(\gamma_0 + \beta) \\ &\quad + \Phi(\beta - \gamma_0) + \beta^2 \Phi(\gamma_0 - \beta) + (1 - \beta^2)\tilde{\Phi}(\gamma_0 + \beta), \end{aligned} \tag{13}$$

$$\begin{aligned} R(\beta, \hat{\beta}_{l_1}) &= \int_{\gamma_1 - \beta}^{\infty} (z - \gamma_1)^2 \phi(z)\, \mathrm{d}z + \int_{-\infty}^{-\gamma_1 - \beta} (z + \gamma_1)^2 \phi(z)\, \mathrm{d}z \\ &\quad + \beta^2 \left[\Phi(\gamma_1 - \beta) - \tilde{\Phi}(\gamma_1 + \beta) \right] \\ &= (-\gamma_1 - \beta)\phi(\gamma_1 - \beta) + (\beta - \gamma_1)\phi(\gamma_1 + \beta) \\ &\quad + (\gamma_1^2 + 1)\Phi(\beta - \gamma_1) + \beta^2 \Phi(\gamma_1 - \beta) + (\gamma_1^2 + 1 - \beta^2)\tilde{\Phi}(\gamma_1 + \beta) \end{aligned} \tag{14}$$

We list a few Gaussian tail bounds here that we will use in the proofs later. Detailed discussion can be found in related articles (Feller, 1968; Donoho and Johnstone, 1994; Foster and George, 1994; Abramovich et al., 2006).

Lemma 1. *For any $z > o$,*

1. $\phi(z)(z^{-1} - z^{-3}) \le \tilde{\Phi}(z) \le \phi(z)z^{-1}$;
2. $\tilde{\Phi}(z) \le e^{-z^2}/2$.
3. $\phi(z)(x^{-1} - x^{-3} + (1\cdot 3)\cdot x^{-5} - (1\cdot 3\cdot 5)\cdot x^{-7} + \cdots + (-1)^k \cdot (2k-1)!! \cdot x^{-2k-1}$ *overestimates* $\tilde{\Phi}(z)$ *if k is even, and underestimates* $\tilde{\Phi}(z)$ *if k is odd.*
4. $\phi(z)(z^{-1} - z^{-3} + (1\cdot 3)\cdot z^{-5} - (1\cdot 3\cdot 5)\cdot z^{-7} + \cdots + (-1)^k \cdot (2k-1)!! \cdot z^{-2k-1}$ *overestimates* $\tilde{\Phi}(z)$ *if k is even, and underestimates* $\tilde{\Phi}(z)$ *if k is odd.*

Lemma 2. *For large enough $\gamma_0 > 0$,*

$$\inf_{\gamma_1} \sup_{\beta} \frac{R(\beta, \hat{\beta}_{l_1})}{R(\beta, \hat{\beta}_{l_0})} > \gamma_0. \tag{15}$$

Proof. It suffices to show that for any fixed γ_0 and any γ_1

$$\sup_{\beta} \frac{R(\beta, \hat{\beta}_{l_1})}{R(\beta, \hat{\beta}_{l_0})} > \gamma_0.$$

Suppose $\gamma_1 \ge \gamma_0/\sqrt{2}$, let $\beta_n = (n+1)\gamma_0$, then

$$\|\hat{\beta}_{l_1} - \hat{\beta}_{LS}\|_2^2 > \|\hat{\beta}_{l_1} - \hat{\beta}_{LS}\|_2^2 I_{\{\hat{\beta}_{LS} > \gamma_1\}} = \gamma_1^2 I_{\{\hat{\beta}_{LS} > \gamma_1\}} \ge \frac{\gamma_0^2}{2} I_{\{\hat{\beta}_{LS} > \gamma_1\}}.$$

Hence,

$$\mathbb{E}\|\hat{\beta}_{l_1} - \hat{\beta}_{LS}\|_2^2 > \frac{\gamma_0^2}{2} P(\hat{\beta}_{LS} > \gamma_1),$$

Thus,

$$\begin{aligned}
\mathbb{E}\|\hat{\beta}_{l_1} - \beta_n\|_2^2 &\ge E\|\hat{\beta}_{l_1} - \hat{\beta}_{LS}\|_2^2 - E\|\hat{\beta}_{LS} - \beta_n\|_2^2 > \frac{\gamma_0^2}{2} P(\hat{\beta}_{LS} > \gamma_1) - 1 \\
&= \left(\frac{\gamma_0^2}{2} - 1\right) P(\hat{\beta}_{LS} > \gamma_1) - P(\hat{\beta}_{LS} \le \gamma_1) \\
&> \gamma_0 \Phi((n+1)\gamma_0 - \gamma_1) - \Phi(\gamma_1 - (n+1)\gamma_0),
\end{aligned}$$

for large enough γ_0 and $Z \sim N(0,1)$.

On the other hand,

$$\begin{aligned}
\mathbb{E}\|\hat{\beta}_{l_0} - \beta_n\|_2^2 &= -n\gamma_0\phi(n\gamma_0) + (n+2)\gamma_0\phi\left((n+2)\gamma_0\right) + \Phi(n\gamma_0) \\
&\quad + (n+1)^2\gamma_0^2\tilde{\Phi}(n\gamma_0) + (1 - (n+1)^2\gamma_0^2)\tilde{\Phi}((n+2)\gamma_0) \\
&\le 1 + \left(-n\gamma_0 - \frac{1}{n\gamma_0} + \frac{(n+1)^2\gamma_0}{n}\right)\phi(n\gamma_0) \\
&\quad + \left((n+2)\gamma_0 + \frac{1 - (n+1)^2\gamma_0^2}{(n+2)\gamma_0}\right)\phi((n+2)\gamma_0) \\
&\le 1 + \left(2 + \frac{1}{n} + 2e^{-2(n+1)\gamma_0^2}\right)\gamma_0\phi(n\gamma_0).
\end{aligned}$$

Hence,

$$\frac{R(\beta_n, \hat{\beta}_{l_1})}{R(\beta_n, \hat{\beta}_{l_0})} \ge \frac{\gamma_0\Phi((n+1)\gamma_0 - \gamma_1) - \Phi(\gamma_1 - (n+1)\gamma_0)}{1 + \left(2 + n^{-1} + 2e^{-2(n+1)\gamma_0^2}\right)\gamma_0\phi(n\gamma_0)}$$

Let $n \to \infty$, then

$$\sup_{\beta} \frac{R(\beta, \hat{\beta}_{l_1})}{R(\beta, \hat{\beta}_{l_0})} \ge \lim_{n\to\infty} \frac{R(\beta_n, \hat{\beta}_{l_1})}{R(\beta_n, \hat{\beta}_{l_0})} \ge \gamma_0. \tag{16}$$

For those $0 \le \gamma_1 < \gamma_0/\sqrt{2}$, we consider $\beta = 0$ and denote

$$R_0(\gamma_0) = R(0, \hat{\beta}_{l_0}(\gamma_0)) = 2\gamma_0\phi(\gamma_0) + 2\tilde{\Phi}(\gamma_0) \tag{17}$$

$$R_1(\gamma_1) = R(0, \hat{\beta}_{l_1}(\gamma_1)) = -2\gamma_1\phi(\gamma_1) + 2(\gamma_1^2 + 1)\tilde{\Phi}(\gamma_1). \tag{18}$$

Consider $0 < c \le \gamma_1 < \gamma_0/\sqrt{2}$. By Lemma 1, for any $z \ge c$, $\tilde{\Phi}(z) - \phi(z)(1/z - 1/z^3 + 1/z^5) \ge 0$. We have $\phi(\gamma_1) \ge \phi(\gamma_0)e^{\gamma_0^2/4}$, and

$$\begin{aligned} R_1(\gamma_1) &\ge -2\gamma_1\phi(\gamma_1) + 2(\gamma_1^2+1)\left(\gamma_1^{-1} - \gamma_1^{-3} + \gamma_1^{-5}\right)\phi(\gamma_1) \\ &\ge 2^{7/2}\gamma_0^{-5}\phi(\gamma_0)e^{\gamma_0^2/4} \\ R_0(\gamma_0) &\le 2\left(\gamma_0 + \gamma_0^{-1}\right)\phi(\gamma_0). \end{aligned}$$

Thus for large enough γ_0

$$\frac{R_1(\gamma_1)}{R_0(\gamma_0)} \ge \frac{2^{5/2}e^{\gamma_0^2/4}}{\gamma_0^6 + \gamma_0^4} > \gamma_0.$$

Lastly, since

$$\frac{d}{d\gamma_1}R_1(\gamma_1) = -4\phi(\gamma_1) + 4\gamma_1\tilde{\Phi}(\gamma_1) < 0,$$

for any $0 \le \gamma_1 \le c$, we have

$$\frac{R_1(\gamma_1)}{R_0(\gamma_0)} \ge \frac{R_1(c)}{R_0(\gamma_0)} > \gamma_0. \qquad \square$$

Proof of Theorem 1. Let $C_1 = \min_{\gamma_0>0}\left\{\frac{2^{5/2}e^{\gamma_0^2/4}}{\gamma_0^6+\gamma_0^4} - \gamma_0\right\} < 5.161$, which occurs at $\gamma_0 = 5.71$. $\square$

Lemma 3. *There exists an $M > 0$ and a constant $C > 0$, such that for all $\gamma_1 > M$,*

$$\inf_{\gamma_0}\sup_{\beta}\frac{R(\beta,\hat{\beta}_{l_0})}{R(\beta,\hat{\beta}_{l_1})} \le 1 + C\gamma_1^{-1} \tag{19}$$

It suffices to show that for all $\beta \ge 0$ and particular values of γ_0, we have

$$\frac{R(\beta,\hat{\beta}_{l_0})}{R(\beta,\hat{\beta}_{l_1})} \le 1 + C\gamma_1^{-1}. \tag{20}$$

The proof is done by generating bounds for the risks at various β's. Before giving these proofs, we need to relate $R(\beta,\hat{\beta}_{l_0})$ to $R(\beta,\hat{\beta}_{l_1})$.

$$\begin{aligned} R(\beta,\hat{\beta}_{l_0}) &= (\gamma_0-\beta)\phi(\gamma_0-\beta) + (\gamma_0+\beta)\phi(\gamma_0+\beta) \\ &\quad +\Phi(-\gamma_0+\beta) + \beta^2\Phi(\gamma_0-\beta) + (1-\beta^2)\tilde{\Phi}(\gamma_0+\beta) \\ &= (\gamma_1+\Delta\gamma-\beta)\phi(\gamma_1-\beta) + (\gamma_1+\Delta\gamma-\beta)\left.\frac{\partial}{\partial\gamma}\phi(\gamma-\beta)\right|_{\gamma_1}\Delta\gamma \\ &\quad +(\gamma_1+\Delta\gamma+\beta)\phi(\gamma_1+\beta) + (\gamma_1+\Delta\gamma+\beta)\left.\frac{\partial}{\partial\gamma}\phi(\gamma+\beta)\right|_{\gamma_1}\Delta\gamma \\ &\quad +\Phi(-\gamma_1+\beta) + \left.\frac{\partial}{\partial\gamma}\Phi(-\gamma+\beta)\right|_{\gamma_1}\Delta\gamma + \beta^2\Phi(\gamma_1-\beta) + \beta^2\left.\frac{\partial}{\partial\gamma}\Phi(\gamma-\beta)\right|_{\gamma_1}\Delta\gamma \\ &\quad +(1-\beta^2)\tilde{\Phi}(\gamma_1+\beta) + (1-\beta^2)\left.\frac{\partial}{\partial\gamma}\tilde{\Phi}(\gamma+\beta)\right|_{\gamma_1}\Delta\gamma + \gamma_1 e^{-\gamma_1^2/2}o(\Delta\gamma) \\ &= (\gamma_1-\beta)\phi(\gamma_1-\beta) + (\gamma_1+\beta)\phi(\gamma_1+\beta) + \Phi(-\gamma_1+\beta) + \beta^2\Phi(\gamma_1-\beta) \\ &\quad +(1-\beta^2)\tilde{\Phi}(\gamma_1+\beta) - (\gamma_1^2-2\beta\gamma_1)\phi(\gamma_1-\beta)\Delta\gamma \\ &\quad -(\gamma_1^2+2\beta\gamma_1)\phi(\gamma_1+\beta)\Delta\gamma + \gamma_1 e^{-\gamma_1^2}o(\Delta\gamma) \\ &= R(\beta,\hat{\beta}_{l_1}) + 2\gamma_1\phi(\gamma_1-\beta) + 2\gamma_1\phi(\gamma_1+\beta) - \gamma_1^2\Phi(-\gamma_1+\beta) - \gamma_1^2\tilde{\Phi}(\gamma_1+\beta) \\ &\quad -(\gamma_1^2-2\beta\gamma_1)\phi(\gamma_1-\beta)\Delta\gamma - (\gamma_1^2+2\beta\gamma_1)\phi(\gamma_1+\beta)\Delta\gamma + \gamma_1 e^{-\gamma_1^2/2}o(\Delta\gamma) \end{aligned}$$

We can now provide separate proofs for β within the following regions:

1. $0 \le \beta \le \gamma_1 - \sqrt{\log(\gamma_1/2)}$

2. $\gamma_1 - \sqrt{\log(\gamma_1/2)} < \beta \leq \gamma_1 + \sqrt{2\log(\gamma_1)}$
3. $\gamma_1 + \sqrt{2\log(\gamma_1)} < \beta$

Proof for case 1, $0 \leq \beta \leq \gamma_1 - \sqrt{\log(\gamma_1/2)}$. Use the trivial estimator $\hat{\beta}_{l_0} = 0$, ie. set $\gamma_0 = \infty$. Then,

$$
\begin{aligned}
R(\beta, \hat{\beta}_0) &= \beta^2 \\
&\quad \text{and} \\
R(\beta, \hat{\beta}_1) &= \mathbb{E}_\beta \left[(Z - \gamma_1)^2 I_{\{\beta + Z > \gamma_1\}} + (Z + \gamma_1)^2 I_{\{\beta + Z < -\gamma_1\}} \right] + (\|\mathbf{x}\|\beta)^2 \mathbb{P}(|\beta + Z| \leq \gamma_1) \\
&> \beta^2 \mathbb{P}(-\gamma_1 - \beta \leq Z \leq \gamma_1 - \beta) \\
&> \beta^2 \mathbb{P}\left(-2\gamma_1 + \sqrt{\log(\gamma_1/2)} \leq Z \leq \sqrt{\log(\gamma_1/2)} \right) \\
&= \beta^2 \left(1 - \Phi\left(-2\gamma_1 + \sqrt{\log(\gamma_1/2)} \right) - \tilde{\Phi}\left(\sqrt{\log(\gamma_1/2)} \right) \right) \\
&= \beta^2 \left(1 - \tilde{\Phi}\left(2\gamma_1 - \sqrt{\log(\gamma_1/2)} \right) - \tilde{\Phi}\left(\sqrt{\log(\gamma_1/2)} \right) \right) \\
&= \beta^2 \left(1 - 2\tilde{\Phi}\left(\sqrt{\log(\gamma_1/2)} \right) \right) \\
&> \beta^2 \left(1 - \exp\left\{ - \left(\sqrt{\log(\gamma_1/2)} \right)^2 \right\} \right) \\
&= \beta^2 \left(1 - \frac{2}{\gamma_1} \right) \\
&\quad \text{therefore} \\
\frac{R(\beta, \hat{\beta}_{l_0})}{R(\beta, \hat{\beta}_{l_1})} &\leq \frac{\beta^2}{\beta^2 \left(1 - \frac{2}{\gamma_1} \right)} \\
&= \frac{\gamma_1}{\gamma_1 - 2} \\
&= 1 + \frac{2}{\gamma_1 - 2} \\
&= 1 + o(\gamma_1^{-1}), \text{ if } \gamma_1 > 2.
\end{aligned}
$$

We have implicitly defined $\frac{0}{0} = 1$, which can be justified in this case using a limit argument. □

Proof for case 2, $\gamma_1 - \sqrt{\log(\gamma_1/2)} < \beta \leq \gamma_1 + \sqrt{2\log(\gamma_1)}$. Recall that

$$
\begin{aligned}
R(\beta, \hat{\beta}_0) - R(\beta, \hat{\beta}_1) = &(2\gamma_1 + 2\Delta\gamma + 2\beta\gamma_1\Delta\gamma - \gamma_1^2\Delta\gamma)\phi(\gamma_1 - \beta) - \gamma_1^2\Phi(-\gamma_1 + \beta) \\
&+ (2\gamma_1 + 2\Delta\gamma - \gamma_1^2\Delta\gamma - 2\beta\gamma_1\Delta\gamma)\phi(\gamma_1 + \beta) - \gamma_1^2\tilde{\Phi}(\gamma_1 + \beta) + o(\Delta\gamma)
\end{aligned}
$$

We want to replace $\phi(\gamma_1 + \beta)$ by $\phi(\gamma_1 - \beta)$. Need to make sure that the term on $\phi(\gamma_1 + \beta)$ is positive, so that this in fact increases the difference. This holds for $\beta \leq \gamma_1 + \sqrt{2\log\gamma_1}$ and $\Delta\gamma = \frac{1}{2\gamma_1}$,

$$
\begin{aligned}
2\gamma_1 + 2\Delta\gamma - \gamma_1^2\Delta\gamma - 2\beta\gamma_1\Delta\gamma &= 2\gamma_1 + 2\frac{1}{2\gamma_1} - \gamma_1^2\frac{1}{2\gamma_1} - 2\beta\gamma_1\frac{1}{2\gamma_1} \\
&= 2\gamma_1 + \frac{1}{\gamma_1} - \frac{\gamma_1}{2} - \beta \\
&> 2\gamma_1 + \frac{1}{\gamma_1} - \frac{\gamma_1}{2} - \gamma_1 - \sqrt{2\log(\gamma_1)} \\
&= \frac{\gamma_1}{2} + \frac{1}{\gamma_1} - \sqrt{2\log(\gamma_1)} \\
&> 0, \text{ if } \gamma_1 \geq 1.
\end{aligned}
$$

Making the replacement above,

$$
\begin{aligned}
R(\beta, \hat{\beta}_0) - R(\beta, \hat{\beta}_1) &< (4\gamma_1 + 4\Delta\gamma - 2\gamma_1^2\Delta\gamma)\phi(\gamma_1 - \beta) + o(\gamma_1^{-1}) \\
&= (4\gamma_1 + 4\frac{1}{2\gamma_1} - 2\gamma_1^2\frac{1}{2\gamma_1})\phi(\gamma_1 - \beta) + o(\gamma_1^{-1}) \\
&< \left(3\gamma_1 + \frac{2}{\gamma_1} \right) \phi(\gamma_1 - \beta) + o(\gamma_1^{-1})
\end{aligned}
$$

Next, we need a lower bound R_{l_1}.

$$\begin{aligned}
R(\beta,\hat{\beta}_{l_1}) &= (-\gamma_1-\beta)\phi(\gamma_1-\beta)+(\beta-\gamma_1)\phi(\gamma_1+\beta)\\
&\quad +(\gamma_1^2+1)\Phi(\beta-\gamma_1)+\beta^2\Phi(\gamma_1-\beta)+(\gamma_1^2+1-\beta^2)\tilde{\Phi}(\gamma_1+\beta)\\
&= (-\gamma_1-\beta)\phi(\gamma_1-\beta)+(\beta-\gamma_1)\phi(\gamma_1+\beta)\\
&\quad +(\gamma_1^2+1)\tilde{\Phi}(\gamma_1-\beta)+\beta^2\Phi(\gamma_1-\beta)+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)\\
&= (-\gamma_1-\beta)\phi(\gamma_1-\beta)+(\beta-\gamma_1)\phi(\gamma_1+\beta)\\
&\quad +(\gamma_1^2+1)(1-\Phi(\gamma_1-\beta))+\beta^2\Phi(\gamma_1-\beta)+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)\\
&= (-\gamma_1-\beta)\phi(\gamma_1-\beta)+(\beta-\gamma_1)\phi(\gamma_1+\beta)\\
&\quad +(\gamma_1^2+1)-(\gamma_1^2+1)\Phi(\gamma_1-\beta))+\beta^2\Phi(\gamma_1-\beta)+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)\\
&= (-\gamma_1-\beta)\phi(\gamma_1-\beta)+(\beta-\gamma_1)\phi(\gamma_1+\beta)\\
&\quad +(\gamma_1^2+1)-(\gamma_1^2+1-\beta^2)\Phi(\gamma_1-\beta))+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)
\end{aligned}$$

For replacement while keeping bounds, separate into two cases: 1) $\beta^2 \geq \gamma_1^2+1$, and 2) $\beta^2 < \gamma_1^2+1$.

$$\begin{aligned}
R(\beta,\hat{\beta}_{l_1}) &= (-\gamma_1-\beta)\phi(\gamma_1-\beta)+(\beta-\gamma_1)\phi(\gamma_1+\beta)\\
&\quad +(\gamma_1^2+1)-(\gamma_1^2+1-\beta^2)\Phi(\gamma_1-\beta))+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)\\
&\quad \text{Case 1}\\
&> -\gamma_1\left(\phi(\gamma_1-\beta)+\phi(\gamma_1+\beta)\right)-\beta\left(\phi(\gamma_1-\beta)-\phi(\gamma_1+\beta)\right)\\
&\quad (\gamma_1^2+1)-(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)\\
&> -\gamma_1-\beta\phi(\gamma-\beta)+\gamma_1^2+1\\
&> -\gamma_1-\gamma_1/2+\gamma_1^2+1\\
&= \gamma_1^2-\frac{3\gamma_1}{2}+1\\
&> 0, \text{ if } \gamma_1>1.\\
&\quad \text{Case 2}\\
&= -\gamma_1\left(\phi(\gamma_1-\beta)+\phi(\gamma_1+\beta)\right)-\beta\left(\phi(\gamma_1-\beta)-\phi(\gamma_1+\beta)\right)\\
&\quad (\gamma_1^2+1)-(\gamma_1^2+1-\beta^2)\Phi(\gamma_1-\beta)+(\gamma_1^2+1-\beta^2)\Phi(-\gamma_1-\beta)\\
&> -\gamma_1-\beta\phi(\gamma-\beta)+(\gamma_1^2+1)-(\gamma_1^2+1-\beta^2)\Phi(\gamma_1-\beta)\\
&> -\frac{3\gamma_1}{2}+\gamma_1^2+1-\gamma_1^2-1+\beta^2\\
&> -\frac{3\gamma_1}{2}+\left(\gamma_1-\sqrt{\log(\gamma_1/2)}\right)^2\\
&> 0, \text{ if } \gamma_1\geq 2.
\end{aligned}$$

Using worst case $\beta^2 < \gamma_1^2+1$, the above yields

$$\begin{aligned}
\frac{R(\beta,\hat{\beta}_{l_0})}{R(\beta,\hat{\beta}_{l_1})} &\leq 1+\frac{\frac{3\gamma_1}{2}+\frac{1}{\gamma_1}+o(\gamma_1^{-1})}{-\frac{3\gamma_1}{2}+\left(\gamma_1-\sqrt{\log(\gamma_1/2)}\right)^2}\\
&= 1+o(\gamma_1^{-1}), \text{ for } \gamma_1\geq 2.
\end{aligned}$$

□

Proof for case 3, $\beta > \gamma_1+\sqrt{2\log(\gamma_1)}$. Let $\Delta\gamma = 0$,

$$\begin{aligned}
R(\beta,\hat{\beta}_{l_0})-R(\beta,\hat{\beta}_{l_1}) &= 2\gamma_1\phi(\gamma_1-\beta)+2\gamma_1\phi(\gamma_1+\beta)-\gamma_1^2\Phi(-\gamma_1+\beta)-\gamma_1^2\tilde{\Phi}(\gamma_1+\beta)\\
&< 4\gamma_1\phi(\sqrt{2\log\gamma_1})-\gamma_1^2\Phi(\sqrt{2\log\gamma_1})\\
&= \frac{4}{\sqrt{2\pi}}-\gamma_1^2\Phi(\sqrt{2\log\gamma_1})\\
&< 0, \text{ if } \gamma_1\geq 1.415.
\end{aligned}$$

□

Proof of Lemma 3. Using the above proofs, let M = 2 and C the constant suppressed in $o(\gamma_1^{-1})$). □

Proof of Theorem 2. For $\gamma_1 < M$ we know that there exists some $\epsilon > 0$ such that $R(\beta,\hat{\beta}_{l_1}(\gamma_1)) \geq \epsilon$ for all β. If we use the trivial estimator $\gamma_0 = 0$, we know it has risk 1. Hence, we can pick $C_2 = \max(1/\epsilon, C)$ where C is from our lemma, then Theorem 2 follows. □

References

Abramovich, F., Benjamini, Y., Donoho, D. L., and Johnstone, I. M. (2006). Adapting to unknown sparsity by controlling the false discovery rate. *Ann. Stat.*, 34(2):584–653.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

Candes, E. and Tao, T. (2007). The Dantzig Selector: Statistical Estimation When p Is Much Larger than n. *The Annals of Statistics*, 35(6):pp. 2313–2351.

Candes, E. J., Wakin, M. B., and Boyd, S. P. (2007). Enhancing Sparsity by Reweighted L1 Minimization.

Donoho, D. L. and Johnstone, J. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455.

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least Angle Regression. *The Annals of Statistics*, 32(2):407–451.

Feller, W. (1968). volume 1. John Wiley & Sons, Inc.

Foster, D. P. and George, E. I. (1994). The Risk Inflation Criterion for Multiple Regression. *The Annals of Statistics*, 22(4):pp. 1947–1975.

Foster, D. P. and Stine, R. A. (2004). Variable selection in data mining: Building a predictive model for bankruptcy. *Journal of the American Statistical Association*, pages 303–313.

Foster, D. P. and Stine, R. A. (2005). Polyshrink: An Adaptive Variable Selection Procedure That Is Competitive with Bayes Experts.

George, E. I. and Foster, D. P. (2000). Calibration and Empirical Bayes Variable Selection. *Biometrika*, 87:731–747.

Johnstone, I. M. and Silverman, B. W. (2005). Empirical Bayes selection of wavelet thresholds. *Annals of Statistics 2005, Vol. 33, No. 4, 1700-1752.*

Meinshausen, N. (2007). Relaxed Lasso. *Computational Statistics & Data Analysis*, 52(1):374–393.

Natarajan, B. K. (1995). Sparse Approximate Solutions to Linear Systems. *SIAM J. Comput.*, 24(2):227–234.

Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464.

Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288.

Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108.

Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67.

Zhou, J., Foster, D. P., Stine, R. A., and Ungar, L. H. (2006). Streamwise Feature Selection. *Journal of Machine Learning Research*, 7:1861–1885.

Zou, H. (2006). The Adaptive Lasso and Its Oracle Properties. *Journal of the American Statistical Association*, 101:1418–1429.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the Elastic Net. *Journal of the Royal Statistical Society, Series B*, 67:301–320.